

Bayesian Gaussian Copula Graphical Modeling for Dupuytren Disease

Abdolreza Mohammadi
University of Groningen
a.mohammadi@rug.nl

Fentaw Abegaz
University of Groningen
f.abegaz.yazew@rug.nl

Edwin van den Heuvel
Eindhoven University of Technology
e.r.v.d.heuvel@tue.nl

Ernst C. Wit
University of Groningen
e.c.wit@rug.nl

September 16, 2015

Abstract

Dupuytren disease is a fibroproliferative disorder with unknown etiology that often progresses and eventually can cause permanent contractures of the affected fingers. In this paper, we provide a computationally efficient Bayesian framework to discover potential risk factors and investigate which fingers are jointly affected. Our Bayesian approach is based on Gaussian copula graphical models, which are one potential way to discover the underlying conditional independence structure of variables in multivariate mixed data. In particular, we combine the semiparametric Gaussian copula with extended rank likelihood which is appropriate to analyse multivariate mixed data with arbitrary marginal distributions. For the graph structure learning, we construct a computationally efficient search algorithm which is a trans-dimensional MCMC algorithm based on a birth-death process. In addition, to make our statistical method easily accessible to other researchers, we have implemented our method in C++ and interfaced with R software as an R package **BDgraph** which is available at <http://CRAN.R-project.org/package=BDgraph>.

Keywords: Dupuytren disease; Risk factors; Bayesian inference; Gaussian copula graphical models; Bayesian model selection; Latent variable models; Birth-death process; Markov chain Monte Carlo.

1 Introduction

Dupuytren disease is a hereditary disorder that is present worldwide. It is however more prevalent in people with northern European ancestry (Bayat and McGrouther, 2006). The disease is an incurable fibroproliferative disorder that alters the palmar fascia of the hand and may causes progressive and permanent flexion contracture of the fingers. Initially, skin pittings and subcutaneous nodules appear in the palm; see Figure 1(a). At a later stage, cords appear that connect the nodules and may contract the fingers into a flexed position; see Figure 1(b). Contracture can arise in a single ray or in multiple rays. The

disease mostly appears on the ulnar side of the hand, i.e., it affects the pink and ring fingers most-frequently (see Figure 6). The only available treatment is surgical intervention. Although much is known about the disease, the questions arising are: (1) What variables affect the disease and in what way? (2) Should surgical intervention focus on single or on multiple fingers? The first is an epidemiological question, the second is a clinical one.

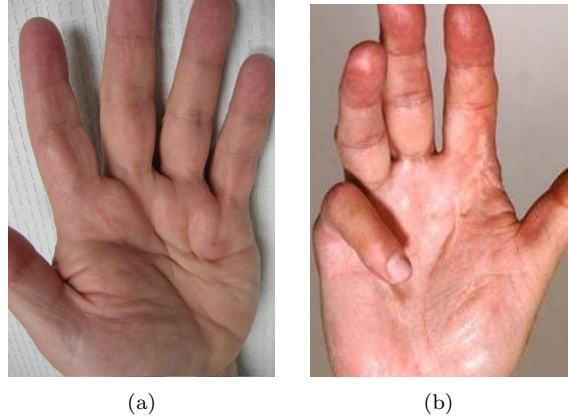


Figure 1: (a) hand image of a patient with Dupuytren disease whose fingers have not been affected by the disease. Palmar nodules and small cords have no signs of contracture. (b) hand image of a patient with Dupuytren disease whose finger have been affected by the disease.

Empirical research has described the patterns of occurrence of Dupuytren disease in multiple fingers. Meyerding (1936) have stated that most often the combination of affected ring and little fingers occurred, followed by the combination of an affected third, fourth, and fifth finger. Tubiana et al. (1982) found that Dupuytren disease rarely affects radial side alone, and that the radial affect is often associate with an affected ulnar side. Milner (2003) noticed that patients who had required surgery because of an affected thumb were on average 8 years older and had suffered significantly longer from the disease compared to patients with a mildly affected radial side. Moreover, these patients suffered from ulnar disease that repeatedly had required surgery, suggesting an intractable form of disease. More recently, Lanting et al. (2014), with a multivariate ordinal logit model, suggested that the middle finger is substantially correlated with other fingers on the ulnar side, and the thumb and index finger are correlated. They took into account age and sex, and tested hypotheses on independence between groups of fingers. However, so far, no serious multivariate analysis of the disease has been performed taking into account potential risk factors.

Essential risk factors of Dupuytren disease include both phenotypic and genotypic factors (Shih and Bayat, 2010), such as genetic predisposition and ethnicity, as well as sex and age. However, it is unclear whether Dupuytren disease is a complex oligogenic or a simple monogenic Mendelian disorder. Several life-style risk factors (some considered controversial) include smoking, excessive alcohol consumption, manual work, and hand trauma. Also several diseases, such as diabetes mellitus and epilepsy, are thought to affect the life-style factors risk of incurring Dupuytren disease. However, the role of these life-style factors

and diseases has not been fully elucidated, and the results of different studies are occasionally conflicting (Lanting et al., 2014).

In this paper we analyze data collected by the Department of Plastic Surgery of the University Medical Center Groningen in the north of Netherlands from patients who have Dupuytren disease. The data are originally described by Lanting et al. (2013) and Lanting et al. (2014). Both hands of patients are examined for signs of Dupuytren disease. These are tethering of the skin, nodules, cords, and finger contractures in patients with cords. Severity of the disease is measured by the angles on each of the 10 fingers. Recorded potential risk factors include smoking habits, alcohol consumption, whether participants had performed manual labor during a significant part of their life, and whether they had sustained hand injury in the past, including surgery. In addition, information about the presence of Ledderhose diabetes, epilepsy, peyronie, knuckle pad, liver disease, and familial occurrence of Dupuytren disease, defined as a first-degree relative with Dupuytren disease, was collected.

The primary aim of this paper is to model the relationships between the risk factors and disease indicators for Dupuytren disease based on the mixed dataset. We propose a computationally efficient Bayesian statistical method based on Gaussian copula graphical models (GCGMs) for discovering the joint conditional independence structure of binary, ordinal or continuous variables simultaneously. Our Bayesian framework is based on that proposed by Dobra and Lenkoski (2011). In GCGMs, the graph selection procedure is embedded inside a semiparametric framework, using the extended rank likelihood (Hoff, 2007). In this paper we design our proposed Bayesian framework for GCGMs based on a computationally efficient search algorithm, using a trans-dimensional MCMC approach based on a continuous-time birth-death process (Mohammadi and Wit, 2015b).

In this work, we follow a similar approach as that of Mohammadi and Wit (2015b) by using their birth-death MCMC algorithm. Changes have been made to allow more general data structures of mixed type using copula graphical models. Importantly we have made substantial improvements in their algorithm to overcome the computational bottle-neck in evaluating normalizing constants. Furthermore, our proposed approach can handle missing data without any additional computational effort, if the missingness is completely at random (MCAR).

In Section 2, we illustrate our Bayesian framework based on Gaussian copula graphical models. In addition, we show the performance of our method and we compare it to state-of-the-art alternatives. In section 3, we analyze the Dupuytren disease dataset based on the proposed Bayesian birth-death MCMC method. In this section, first analyze potential phenotype risk factors for Dupuytren disease. Moreover, we analyze the relationship between consider the severity of Dupuytren disease of pairs of fingers. The result helps surgeons to decide whether they should operate one finger or they should operate multiple fingers simultaneously. Finally, we discuss the connections between existing methods and possible future directions.

2 Methodology

Graphical models (Lauritzen, 1996) provide an effective way to describe statistical patterns in multivariate data. In this context undirected Gaussian graphical

models are commonly used, since inference in such models is often tractable. In undirected Gaussian graphical models, the graph structure is characterized by its precision matrix (the inverse of covariance matrix): the non-zero entries in the precision matrix show the edges in the conditional independence graph. In the real world, data are often non-Gaussian, like the dataset considered in our application. For non-Gaussian continuous data, variables can be transformed to Gaussian latent variables. For discrete data, however, there is no one-to-one transformation into latent Gaussian variables. A common approach is to apply a Markov chain Monte Carlo method (MCMC) to simulate both the latent Gaussian variables and the posterior distributions (Hoff, 2007). Another Bayesian approach is the Gaussian copula graphical models developed by Dobra and Lenkoski (2011), in which the sampler algorithm is based on reversible-jump MCMC and a Cholesky decomposition of the precision matrix. Alternatively, our proposed method implements the birth-death MCMC approach (Mohammadi and Wit, 2015b) which has several computational advantages compared to reversible-jump MCMC approach as we show in our simulation examples.

For studying the multivariate dependence structure of finger contractures and risk factors, we are interested in exploring the conditional independence graph space and identifying conditional independence graphs that are most appropriate for our given data. In this regard, we calculate the posterior distribution of the conditional independence graph G conditional on data

$$Pr(G|\text{data}) = \frac{Pr(G)Pr(\text{data}|G)}{\sum_{G \in \mathcal{G}} Pr(G)Pr(\text{data}|G)},$$

in which \mathcal{G} is the conditional independence graph space. Computing this posterior distribution is computationally unfeasible, since in the denominator we require the sum over all possible graphs. The graph space increases super-exponentially with the dimension of the variables. For p nodes in a conditional independence graph, there are $p(p-1)/2$ possible edges, and hence we have $2^{p(p-1)/2}$ different possible graphs corresponding to all combinations of individual edges being in or out of the graph. For example, in our data we have 23 variables ($p = 23$), resulting in a total number of possible conditional independence graphs of more than 10^{76} . This motivates us to develop effective search algorithms for exploring graphical model uncertainty. In order to be accurate and scalable, the key is to design computationally efficient search algorithms that are able to quickly move towards high posterior probability regions, and to take advantage of local computations.

2.1 Gaussian copula graphical models

In graphical models, conditional dependence relationships among random variables are presented as a graph G . A graph $G = (V, E)$ specifies a set of vertices $V = \{1, 2, \dots, p\}$, where each vertex corresponds to a random variable, and a set of edges $E \subset V \times V$ (Lauritzen, 1996). \bar{E} denotes the set of non-existing edges. We focus here on undirected graphical models, where $(i, j) \in E \Leftrightarrow (j, i) \in E$, also known as Markov random fields. The absence of an edge between two vertices specifies the pairwise conditional independence of these two variables given the remaining variables, while an edge between two variables determines the conditional dependence of the variables. In our application, for example,

disease risk factors (such as disease factors, alcohol, and hand injury) will be the nodes, and dependencies among them will be the edges.

A graphical model that follows a multivariate normal distribution is called a Gaussian graphical models, also known as a covariance selection model (Dempster, 1972). Zero entries in the precision matrix correspond to the absence of edges on the graph and conditional independence between pairs of random variables given all other variables. We define a zero mean Gaussian graphical model with respect to the graph G as

$$\mathcal{M}_G = \{\mathcal{N}_p(0, \Sigma) \mid K = \Sigma^{-1} \in \mathbb{P}_G\},$$

where \mathbb{P}_G denotes the space of $p \times p$ positive definite matrices with entries (i, j) equal to zero whenever $(i, j) \in \bar{E}$. Let $\mathbf{z} = (\mathbf{z}^1, \dots, \mathbf{z}^n)$ be an independent and identically distributed sample of size n from model \mathcal{M}_G , where \mathbf{z}^i is a p dimensional vector of variables. Then, the likelihood is

$$P(\mathbf{z} | K, G) \propto |K|^{n/2} \exp \left\{ -\frac{1}{2} \text{tr}(KS) \right\}, \quad (1)$$

where $S = \mathbf{z}'\mathbf{z}$.

In line with extending the idea of Gaussian graphical modeling for non-Gaussian data, the Gaussian copula has been considered for graphical modeling among a set of mixed variables (Dobra and Lenkoski, 2011). A copula is a multivariate cumulative distribution function whose uniform marginals are on the interval $[0, 1]$. Copulas provide a flexible tool for understanding dependence among random variables, in particular for non-Gaussian multivariate data, for example, the type of data application we consider in this study which consists of 23 mixed variables (13 phenotype risk factors and 10 variables on the severity of Dupuytren disease in all the hand fingers) of the type discrete, binary, and ordered categorical variables; see Section 3.1.

By Sklar's theorem (Sklar, 1959) there exists a copula C such that any p dimensional distribution function H can be completely specified by its marginal distributions F_j , $j = 1, \dots, p$, and a copula C satisfying

$$H(y_1, \dots, y_p) = C(F_1(y_1), \dots, F_p(y_p)).$$

Conversely, a copula function can be extracted from any p dimension distribution function H and marginal distributions F_j by

$$C(u_1, \dots, u_p) = H(F_1^{-1}(u_1), \dots, F_p^{-1}(u_p)),$$

where $F_j^{-1}(s) = \inf\{t \mid F_j(t) \geq s\}$ are the pseudo-inverse of F_j . The Gaussian copula is given by

$$C(u_1, \dots, u_p \mid \Gamma) = \Phi_p(\Phi^{-1}(u_1), \dots, \Phi^{-1}(u_p) \mid \Theta),$$

where $u_j = F_j(y_j)$, $j = 1, \dots, p$, $\Phi_p(\cdot)$ is the cumulative distribution of a multivariate normal distribution and $\Phi(\cdot)$ is a cumulative distribution function of a univariate normal distribution.

The decomposition of a joint distribution into marginal distributions and a copula suggests that the copula captures the essential dependence features

between random variables. Moreover, the copula measure of dependence is invariant to any monotone transformation of the random variables. Thus, copulas allow one to model the marginal distributions and the dependence structure of a multivariate random variables separately. In copula modeling, Genest et al. (1995) develop a popular semiparametric estimation approach or rank likelihood based estimation in which the association among variables is represented with a parametric copula but the marginals are treated as nuisance parameters. The marginals are estimated non-parametrically using the scaled empirical distribution function $\hat{F}_j(y) = \frac{n}{n+1}F_{n_j}(y)$, where $F_{n_j}(y) = \frac{1}{n} \sum_{i=1}^n I\{y_{ij} \leq y\}$. As a result estimation and inference are robust to misspecification of marginal distributions.

However, for discrete data, where the distribution of ranks depends on the univariate marginal distributions, the use of the semiparametric estimator is somewhat inappropriate for the analysis of mixed continuous and discrete data (Hoff, 2007). To overcome this, he proposes the extended rank likelihood, which is a type of marginal likelihood approach where the ranks are free of the marginal distributions of the discrete data. This makes the extended rank likelihood approach more appropriate for graphical modeling that avoids the difficult problem of modeling the marginal distributions.

Let Y be a collection of continuous, binary, ordinal or count variables with F_j the marginal distribution of Y_j and F_j^{-1} its pseudo inverse. For constructing a joint distribution of Y , we introduce a multivariate normal latent variable as follows

$$Z_1, \dots, Z_n \stackrel{iid}{\sim} \mathcal{N}(0, \Sigma),$$

where Σ is the correlation matrix. We define the observed data as

$$Y_{ij} = F_j^{-1}(\Phi(Z_{ij})).$$

A Gaussian copula-based joint cumulative distribution of \mathbf{Y} is given by

$$P(Y_1 \leq y_1, \dots, Y_p \leq y_p) = \Phi_p(\Phi^{-1}(F_1(y_1)), \dots, \Phi^{-1}(F_p(y_p)) | \Gamma). \quad (2)$$

Our aim is to infer the underlying graph structure G of the mixed variables \mathbf{Y} implied by the continuous latent variables \mathbf{Z} . The observed mixed data from a sample of n observations \mathbf{y} are related to the latent samples \mathbf{z} that belong to the set

$$\mathcal{A}(\mathbf{y}) = \left\{ \mathbf{z} \in \mathbb{R}^{n \times p} : \max_{r, k, s = 1, \dots, n; j = 1, \dots, p} \left\{ z_j^{(k)} : y_j^{(s)} < y_j^{(r)} \right\} < z_j^{(r)} < \min \left\{ z_j^{(s)} : y_j^{(r)} < y_j^{(s)} \right\} \right\}, \quad (3)$$

It follows that inference on the latent space can be performed by substituting the observed data \mathbf{y} with the event $\mathbf{z} \in \mathcal{A}(\mathbf{y})$. For a given graph G and precision matrix $K = \Sigma^{-1}$, the extended rank likelihood is defined as

$$P(\mathbf{z} \in \mathcal{A}(\mathbf{y}) | K, G) = P(\mathbf{z} \in \mathcal{A}(\mathbf{y}) | K, G) = \int_{\mathcal{A}(\mathbf{y})} P(\mathbf{z} | K, G) d\mathbf{z}, \quad (4)$$

where the expression inside the integral for the Gaussian copula based probability function given by (2) takes a similar form as in 1. In the next sections we develop a Bayesian approach based on the extended rank likelihood given in (4).

2.2 Bayesian Gaussian copula graphical models

In the Bayesian framework that follows, we infer about graph and precision matrix (K, G) by considering a posterior distribution

$$P(K, G | \mathbf{z} \in \mathcal{A}(\mathbf{y})) \propto P(\mathbf{z} \in \mathcal{A}(\mathbf{y}) | K, G) P(K | G) P(G), \quad (5)$$

where $P(\mathbf{z} \in \mathcal{A}(\mathbf{y}) | K, G)$ is the extended rank likelihood defined in (4), $P(K | G)$ denotes a prior distribution of a precision matrix K for a given graph structure and $P(G)$ denotes a prior distribution for a graph G .

In particular, we develop a simple and efficient continuous birth-death Markov chain Monte Carlo (MCMC) algorithm for the posterior computation that converges much faster than the reversible MCMC algorithm in Dobra and Lenkoski (2011) and overcomes the computational bottle-neck in Mohammadi and Wit (2015b). Moreover, we evaluate the results induced by the latent variables using posterior predictive analysis on the scale of the original mixed variables.

2.2.1 Prior specification

In what follows we briefly describe the specification of prior distributions for the graph G and the precision matrix K . For the prior distribution of the graph, we propose to use a discrete uniform distribution over the graph space, as a non-informative prior. We also note that other choices of priors for the graph structure have been considered by modeling the joint state of the edges (Scutari, 2013), encouraging sparse graphs (Jones et al., 2005) or a truncated Poisson distribution on the graph size (Mohammadi and Wit, 2015b).

We consider the G-Wishart (Roverato, 2002) distribution as prior distribution of the precision matrix. The G-Wishart is the Wishart distribution restricted to the space of precision matrices with zero entries specified by a graph G , \mathbb{P}_G . The G-Wishart density for $K \in \mathbb{P}_G \sim W_G(b, D)$ can be written as

$$P(K | G) = \frac{1}{I_G(b, D)} |K|^{(b-2)/2} \exp \left\{ -\frac{1}{2} \text{tr}(DK) \right\},$$

where $b > 2$ is the degree of freedom, D is a symmetric positive definite matrix, and $I_G(b, D)$ is the normalizing constant,

$$I_G(b, D) = \int_{\mathbb{P}_G} |K|^{(b-2)/2} \exp \left\{ -\frac{1}{2} \text{tr}(DK) \right\} dK.$$

Dealing with this normalizing constant has been a major issue in recent literature; see Section 2.2.3.

Since the G-Wishart prior is conjugate to normally distributed data (1), the posterior distribution of K condition on graph G is

$$P(K | Z \in \mathcal{A}(\mathbf{y}), G) = \frac{1}{I_G(b^*, D^*)} |K|^{(b^*-2)/2} \exp \left\{ -\frac{1}{2} \text{tr}(D^* K) \right\},$$

where $b^* = b + n$ and $D^* = D + S$ with $S = \mathbf{z}'\mathbf{z}$, that is a G-Wishart distribution, $W_G(b^*, D^*)$. For other choices of priors for the precision matrix see Wang and Pillai (2013), Wang et al. (2015), Wang (2012), Wong et al. (2003).

2.2.2 Sampling algorithm for posterior inference

Sampling from the joint posterior distribution (5) can be done by a computationally efficient birth-death MCMC algorithm proposed in Mohammadi and Wit (2015b) for Gaussian graphical models. Here we extend their algorithm for the more general case of Gaussian copula graphical models. Our algorithm is based on a continuous time birth-death Markov process in which the algorithm explores the graph space by adding or removing an edge in a birth or death event, respectively. The birth and death rates of edges are determined by the stationary distribution of the process. The algorithm is designed in such a way that the stationary distribution equals to the target joint posterior distribution of the graph and the precision matrix (5). The time between two successive events has an exponential distribution. Therefore, the probability of birth and death events are proportional to their rates.

Mohammadi and Wit (2015b, section 3) proved that the birth-death MCMC (BDMCMC) algorithm converges to the target joint posterior distribution of the graph and the precision matrix, by considering the following birth and death rates,

$$\beta_e(K) = \frac{P(G^{+e}, K^{+e} \setminus (k_{ij}, k_{jj}) | Z \in \mathcal{A}(\mathbf{y}))}{P(G, K \setminus k_{jj} | Z \in \mathcal{A}(\mathbf{y}))}, \quad \text{for each } e \in \overline{E}, \quad (6)$$

$$\delta_e(K) = \frac{P(G^{-e}, K^{-e} \setminus k_{jj} | Z \in \mathcal{A}(\mathbf{y}))}{P(G, K \setminus (k_{ij}, k_{jj}) | Z \in \mathcal{A}(\mathbf{y}))}, \quad \text{for each } e \in E, \quad (7)$$

in which $G^{+e} = (V, E \cup \{e\})$ for the inclusion of an edge from the graph G , and $K^{+e} \in \mathbb{P}_{G^{+e}}$ represents an updated precision matrix when an edge is included, and similarly $G^{-e} = (V, E \setminus \{e\})$, and $K^{-e} \in \mathbb{P}_{G^{-e}}$ represent exclusion of an edge and its updated precision matrix. We refer our proposed approach as the extended birth-death MCMC (EBDMCMC) algorithm. Details of the EBDMCMC algorithm are given in Algorithm 1.

In Algorithm 1, the first step is to sample the latent variables given the observed data. Then, based on this sample, we calculate the birth and death rates and waiting times. The birth and death rates are used to calculate the jump type. Details of how to efficiently calculate the birth and death rates are discussed in Section 2.2.3. Finally in step 3, according to the jump, we sample a new precision matrix using a direct sampling scheme from the G-Wishart distribution which is described in Algorithm 2 in Section 2.2.4.

For our algorithm, the Rao-Blackwellized sample mean (Cappé et al., 2003, subsection 2.5) provides an effective way to estimate the posterior probability of each graph. The Rao-Blackwellized estimate of the posterior graph probability is the proportion to the total waiting times for that graph (see Figure 2 in the right). The waiting times for each graph act as the weights of that graph (e.g. $\{W_1, W_2, W_3, \dots\}$ in Figure 2 in the left).

2.2.3 Computing the birth and death rates

Calculating the birth and death rates (6) and (7) has been a major bottle-neck of the EBDMCMC algorithm. Here, we explain how to resolve the computational bottle-neck and come-up with an efficient way to calculate the death rates; the birth rates are calculated in a similar manner.

Algorithm 1 Given a graph $G = (V, E)$ with a precision matrix K , iterate the following steps:

1. Sample the latent data. For each $r \in V$ and $j \in \{1, 2, \dots, n\}$, we update the latent value $z_r^{(j)}$ from its full conditional distribution

$$Z_r|K, Z_{V \setminus \{r\}} = z_{K, V \setminus \{r\}}^{(j)} \sim \mathcal{N}\left(-\sum_{r'} K_{rr'} z_{r'}^{(j)} / K_{rr}, 1/K_{rr}\right), \quad (8)$$

truncated to the interval in (3). This sampling step can be easily modified to handle data that are missing-at-random. That is if y_r is missing, then the full conditional $Z_r|K, Z_{V \setminus \{r\}}$ is the untruncated multivariate normal distribution given in (8).

2. Sample the graph based on birth and death process.

- 2.1. Calculate the birth rates by equation 6 and $\beta(K) = \sum_{e \in \bar{E}} \beta_e(K)$,
- 2.2. Calculate the death rates by equation 7 and $\delta(K) = \sum_{e \in E} \delta_e(K)$,
- 2.3. Calculate the waiting time by $W(K) = 1/(\beta(K) + \delta(K))$,
- 2.4. Calculate the jump type (birth or death),

3. Sample the new precision matrix, according to the jump type, based on Algorithm 2.
-

Following Mohammadi and Wit (2015b) and after some simplification, for each $e = (i, j) \in E$, we have

$$\delta_e(K) = \frac{P(G^{-e})}{P(G)} \frac{I_G(b, D)}{I_{G^{-e}}(b, D)} \left(\frac{D_{jj}^*}{2\pi(k_{ii} - k_{11}^1)}\right)^{\frac{1}{2}} H(K, D^*), \quad (9)$$

where

$$H(K, D^*) = \exp\left\{-\frac{1}{2}\left[\text{tr}(D_{e,e}^*(K^0 - K^1)) - (D_{ii}^* - \frac{(D_{ij}^*)^2}{D_{jj}^*})(k_{ii} - k_{11}^1)\right]\right\},$$

in which

$$K^0 = \begin{bmatrix} k_{ii} & 0 \\ 0 & K_{j, V \setminus j} (K_{V \setminus j, V \setminus j})^{-1} K_{V \setminus j, j} \end{bmatrix},$$

and $K^1 = K_{e, V \setminus e} (K_{V \setminus e, V \setminus e})^{-1} K_{V \setminus e, e}$. The computational bottle-neck in (9) is the ratio of normalizing constants, $\frac{I_G(b, D)}{I_{G^{-e}}(b, D)}$.

Dealing with calculation of normalizing constants. Calculating the ratio of normalizing constants has been a major issue in recent literature (Uhler et al., 2014, Wang and Li, 2012, Mohammadi and Wit, 2015b). To compute the normalizing constants of a G-Wishart, Roverato (2002) proposed an importance sampling algorithm, while Atay-Kayis and Massam (2005) developed a Monte Carlo method. These methods can be computationally expensive and numerically unstable (Jones et al., 2005, Wang and Li, 2012). Wang and Li (2012),

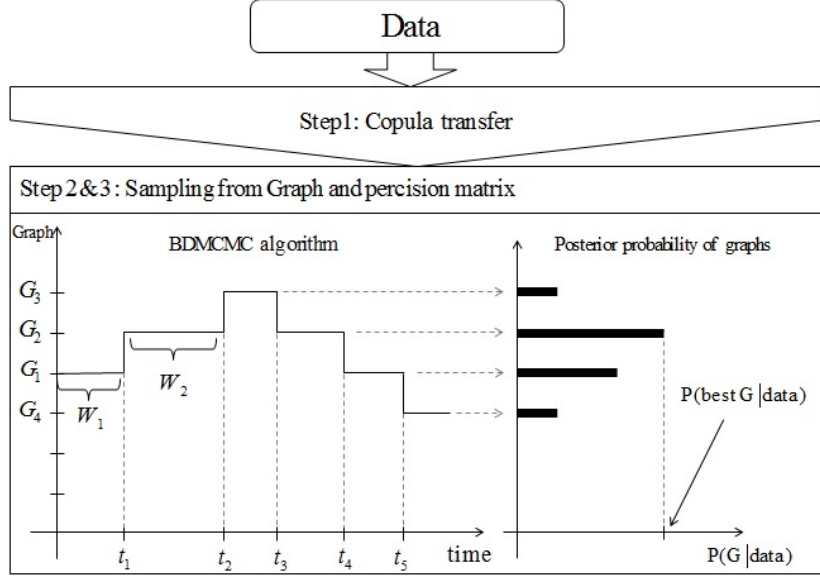


Figure 2: This image visualizes Algorithm 1. (On top) Mixed observed data transformation using the copula to sample the latent variables, (Bottom left) Continuous time EBDMCMC algorithm where $\{W_1, W_2, \dots\}$ denote waiting times and $\{t_1, t_2, \dots\}$ denote jumping times. (Bottom right) Estimated posterior probability of the graphs which are proportional to sum of their waiting times.

Cheng et al. (2012), Mohammadi and Wit (2015b) developed an alternative approach, which borrows ideas from the exchange algorithm (Murray et al., 2012) and the double Metropolis-Hastings (MH) algorithm (Liang, 2010) to compute the ratio of such normalizing constants. However, it has been observed that in case of high dimensional setting (large p), the double MH sampler become computationally inefficient due to the curse of dimensionality. To remedy this, more recently Uhler et al. (2014, theorem 3.7) derived an explicit representation of the normalizing constant ratio.

Theorem 2.1 (Uhler et al. 2014). *let $G = (V, E)$ be an undirected graph and $G^{-e} = (V, E^{-e})$ denotes the graph G with one less edge e . Then*

$$\frac{I_G(b, \mathbb{I}_p)}{I_{G^{-e}}(b, \mathbb{I}_p)} = 2\sqrt{\pi} \frac{\Gamma((b+d+1)/2)}{\Gamma((b+d)/2)},$$

where d denotes the number of triangles formed by the edge e and two other edges in G and \mathbb{I}_p denotes an identity matrix with p dimension.

Therefore, for the case of $D = \mathbb{I}_p$, we have a simplified expression for the death rates in (9), which is given by

$$\delta_e(K) = \frac{P(G^{-e})}{P(G)} \frac{\Gamma((b+d+1)/2)}{\Gamma((b+d)/2)} \left(\frac{2D_{jj}^*}{(k_{ii} - k_{ii}^1)} \right)^{\frac{1}{2}} H(K, D^*).$$

2.2.4 Sampling from posterior distribution of precision matrix

Several sampling methods from a G-Wishart have been proposed for sampling the precision matrix; for a review of existing methods see Wang and Li (2012),

Mitsakakis et al. (2011), Wang et al. (2010). Here we use an exact sampler algorithm developed by Lenkoski (2013) after determining the graph structure and its precision matrix using Algorithm 1 and summarized the details in Algorithm 2.

Algorithm 2. Direct sampler from precision matrix (Lenkoski, 2013). Given a graph $G = (V, E)$ with precision matrix K determined by the jump type (birth for inclusion of an edge and death for exclusion of an edge) from Algorithm 1:

1. Set $\Sigma = K^{-1}$,
 2. Repeat for $j = 1, \dots, p$, until convergence:
 - 2.1 Let $N_j \subset V$ be the set of variables that connected to j in G .
Form Σ_{N_j} and $K_{N_j, j}^{-1}$ and solve

$$\hat{\beta}_j^* = \Sigma_{N_j}^{-1} K_{N_j, j}^{-1},$$
 - 2.2 Form $\hat{\beta}_j \in R^{p-1}$ by plugging zeroes in those locations not connected to j in G and padding the elements of $\hat{\beta}_j^*$ to the rest locations,
 - 2.3 Replace $\Sigma_{j, -j}$ and $\Sigma_{-j, j}$ with $\Sigma_{-j, -j} \hat{\beta}_j$,
 3. Return $K = \Sigma^{-1}$.
-

2.2.5 Simulation study

We perform a comprehensive simulation study with respect to different graph structures to evaluate the performance of the proposed EBDMCMC method and compare it to an alternative approach proposed by Dobra and Lenkoski (Dobra and Lenkoski, 2011), referred to as DL. We generate mixed data from a latent Gaussian copula model with 5 different types of variables, that includes “Gaussian”, “non-Gaussian”, “ordinal”, “count”, and “binary”. We performed all computations with our R package **BDgraph** (Mohammadi and Wit, 2015a).

Corresponding to different sparsity patterns, we consider 3 different kinds of synthetic graphical models, having p nodes:

1. *Random Graph*: A graph in which the edge set E is randomly generated from independent Bernoulli distributions with probability $2/(p-1)$ and corresponding precision matrix is generated from $K \sim W_G(3, I_p)$.
2. *Cluster Graph*: A graph in which the number of clusters is $\max\{2, \lfloor p/20 \rfloor\}$. Each cluster has the same structure as a random graph. The corresponding precision matrix is generated from $K \sim W_G(3, I_p)$.
3. *Scale-free Graph*: A scale-free graph has a power-low degree distribution generated by the Barabasi-Albert algorithm (Albert and Barabási, 2002). The corresponding precision matrix is generated from $K \sim W_G(3, I_p)$.

Figure 3 presents the patterns for the 3 graph structures with $p = 40$ nodes. For each graphical model, we consider different scenarios based on different number

of variables $p = \{10, 20, 30, 40\}$ and different sample size; see Table 1. For each scenario, we generate mixed data and fit our proposed EBDPMC and DL approaches using a uniform prior for the graph, G , and the G-Wishart prior $W_G(3, I_p)$ for the precision matrix, K . We run the two algorithms with the same starting points using 100,000 iterations and 50,000 iterations as a burn-in. Computations for these scenarios were performed in parallel on a 235 batch nodes with 12 cores and 24 GB of memory, running Linux.

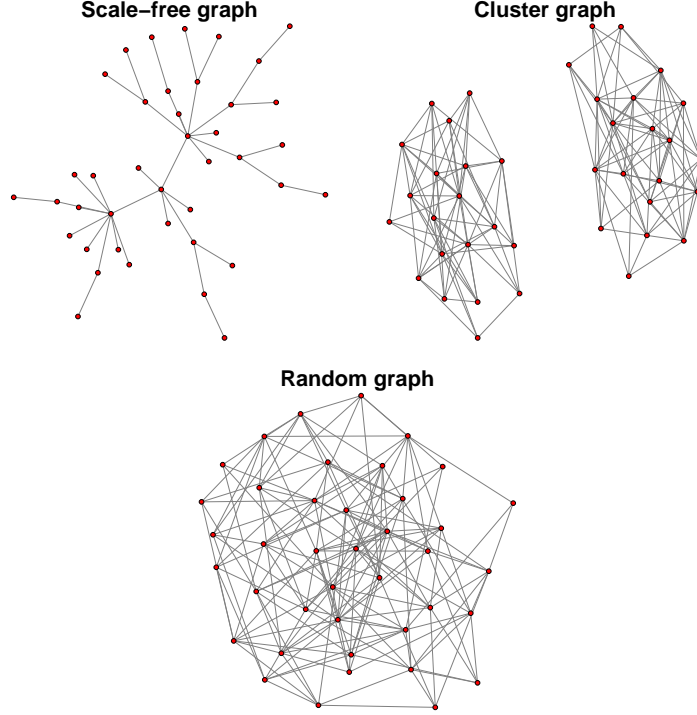


Figure 3: An illustration of the 3 simulated graph structures with $p = 40$ nodes for our simulation example 2.2.5.

To assess the performance of the graph structure, we compute the F_1 -score measure (Powers, 2011) which is defined as

$$F_1\text{-score} = \frac{2\text{TP}}{2\text{TP} + \text{FP} + \text{FN}}, \quad (10)$$

where TP, FP, and FN are the number of true positives, false positives, and false negatives, respectively. The F_1 -score lies between 0 and 1, where 1 stands for perfect identification and 0 for bad identification. Also, we use the mean square error (MSE), defined as

$$\text{MSE} = \sum_e (\hat{p}_e - I(e \in G_{\text{true}}))^2, \quad (11)$$

where \hat{p}_e is the posterior edge inclusion probabilities and $I(e \in G_{\text{true}})$ is an indicator function, such that $I(e \in G_{\text{true}}) = 1$ if $e \in G_{\text{true}}$ and zero otherwise. We calculate the posterior edge inclusion probabilities based on the

Rao-Blackwellization (Cappé et al., 2003, subsection 2.5) for each possible edge $e = (i, j)$ as

$$\hat{p}_e = \frac{\sum_{t=1}^N I(e \in G^{(t)}) W(K^{(t)})}{\sum_{t=1}^N W(K^{(t)})}, \quad (12)$$

where N is the number of iterations and $W(K^{(t)})$ is the waiting time for the graph $G^{(t)}$ with the precision matrix $K^{(t)}$.

p	n	graph	F1-score		MSE	
			EBDMCMC	DL	EBDMCMC	DL
10	30	Random	0.52 (0.16)	0.38 (0.15)	6.44 (2.35)	6.33 (1.8)
		Cluster	0.58 (0.14)	0.43 (0.14)	5.12 (1.75)	5.38 (1.3)
		Scale-free	0.53 (0.18)	0.43 (0.13)	6.46 (2.03)	6.43 (1.31)
10	100	random	0.71 (0.15)	0.67 (0.14)	3.96 (1.73)	4.1 (1.42)
		Cluster	0.68 (0.16)	0.67 (0.16)	3.84 (1.54)	3.49 (1.14)
		Scale-free	0.67 (0.14)	0.63 (0.14)	4.5 (2)	4.16 (1.18)
20	70	random	0.55 (0.08)	0.45 (0.06)	23.28 (5.28)	24.04 (4.00)
		Cluster	0.55 (0.09)	0.47 (0.06)	23.67 (5.36)	21.84 (3.22)
		Scale-free	0.49 (0.13)	0.39 (0.08)	19.97 (6.29)	14.22 (1.96)
20	200	random	0.74 (0.07)	0.61 (0.07)	15.10 (3.63)	17.27 (4.03)
		Cluster	0.73 (0.07)	0.62 (0.07)	14.07 (4.00)	14.58 (3.87)
		Scale-free	0.67 (0.10)	0.55 (0.08)	11.3 (3.68)	9.38 (1.84)
30	100	Random	0.54 (0.06)	0.44 (0.04)	52.3 (9.9)	59.1 (8.7)
		Cluster	0.56 (0.05)	0.47 (0.04)	48.0 (6.5)	54.4 (8.1)
		Scale-free	0.53 (0.17)	0.30 (0.05)	27.7 (14.6)	25.8 (1.7)
30	500	Random	0.79 (0.04)	0.63 (0.07)	25.8 (6.5)	41.1 (14.3)
		Cluster	0.79 (0.05)	0.66 (0.05)	26.3 (5.2)	35.1 (7.9)
		Scale-free	0.81 (0.07)	0.59 (0.06)	9.4 (3.2)	11.7 (3.0)
40	400	Random	0.71 (0.03)	0.57 (0.04)	61.84 (7.77)	81.77 (14.25)
		Cluster	0.71 (0.04)	0.59 (0.04)	58.25 (6.63)	69.4 (13.31)
		Scale-free	0.67 (0.09)	0.46 (0.06)	23.26 (8.37)	23.29 (6.13)
40	800	Random	0.77 (0.03)	0.62 (0.06)	50.07 (8.39)	76.22 (20.43)
		Cluster	0.80 (0.03)	0.68 (0.03)	43.98 (5.93)	58.95 (7.27)
		Scale-free	0.78 (0.07)	0.55 (0.05)	15.05 (5.14)	17.22 (5.4)

Table 1: Summary of performance measures in simulation example 2.2.5 for our method and DL (Dobra and Lenkoski, 2011). The table presents the F_1 -score, defined in (10) and MSE, defined in (11), with 50 replications and standard deviations in parenthesis. The F_1 -score reaches its best score at 1 and its worst at 0. The MSE is positive value for which 0 is minimal and smaller is better. The best models for both F_1 -score and MSE are boldfaced.

Table 1 reports comparisons of the EBDMCMC method with DL, where we

repeat the experiments 50 times and report the average F_1 -score and MSE with their standard errors in parentheses. Our EBDMCMC method performs well overall as its F_1 -score is larger and its MSE is lower than the DL method in most of the cases, mainly because of its fast convergence rate. As we expected, the DL approach converges more slowly compared to our method.

Another way to check the performance of EBDMCMC algorithm and compare it with DL is based on the ROC curves as shown in Figures 4 and 5. The ROC curves depicting the performance of edge selection for each simulated graphs for the case $p = 10$ and $n = 10$, by varying thresholds on the posterior edge inclusion probabilities. As can be seen from the ROC curves the EBDMCMC performs better than the DL method.

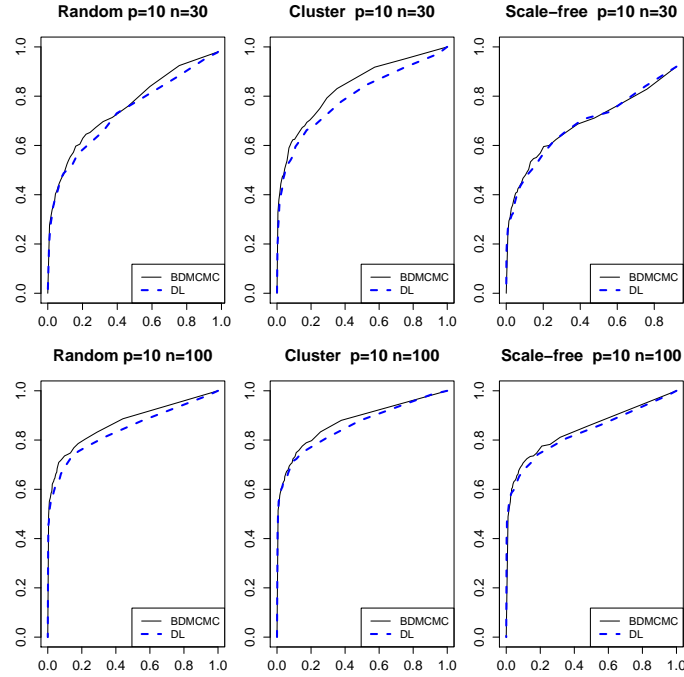


Figure 4: ROC curves depicting the performances of the simulated graphs for the case $p = 10$ and $n = \{30, 100\}$, based on the posterior edge inclusion probabilities.

3 Analysis of Dupuytren disease data

Here we analyze the data collected on patients who have Dupuytren disease in both hands from north of Netherlands by the Department of Plastic Surgery of the University Medical Center Groningen. The data are originally described by Lanting et al. (2013) and Lanting et al. (2014). The data consist of 279 patients who have Dupuytren disease ($n = 279$); among those patients, 79 of them have an irreversible flexion contracture in at least on one of their fingers. Therefore, the data consist of a lot of zeros as shown in Figure 6. That is, though the hands, are affected by the disease, the fingers didn't show any sign of contraction and the total angle measure is taken as 0.

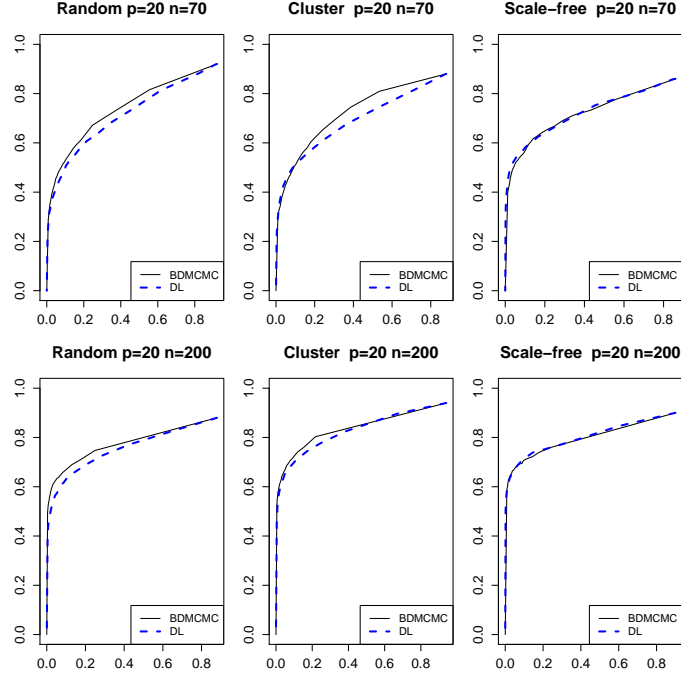


Figure 5: ROC curves depicting the performances of the simulated graphs for the case $p = 20$ and $n = \{70, 200\}$, based on the posterior edge inclusion probabilities.

The severity of the disease in all 10 fingers of the patients is measured by the angles of each finger which is the sum of angles for metacarpophalangeal joints. To study the potential phenotype risk factors of Dupuytren disease, we consider potential phenotype risk 13 factors. These are smoking habits (**Smoking**), alcohol consumption (**Alcohol**), whether participants performed manual labour during a significant part of their life (**Labour**), whether they had sustained hand injury in the past including surgery (**HandInjury**), disease history information about the presence of **Ledderhose**, **Diabetes**, **Epilepsy**, **Peyronie**, **Knuckle pad**, and liver disease (**LiverDiseas**) and familial occurrence of Dupuytren disease which is defined as a first-degree relative with Dupuytren disease (**Relative**).

For each finger we measure angles of metacarpophalangeal joints, two interphalangeal joints (for thumb fingers we only measure two interphalangeal joints); then we sum those angles for each fingers as a measure of the severity of Dupuytren disease. The total angles ideally can vary from 0 to 270 degrees; however, in this dataset the minimum degree is 0 and maximum 157 degrees. The age of participants (in years) ranges from 40 to 89 years, with an average age of 66 years. Smoking is binned into 3 ordered categories (never, stopped, and smoking). Amount of alcohol consumption is binned into 8 ordered categories (ranging from no alcohol to more than 20 consumption per week). All other variables are binary.

In Section 3.1, we infer the Dupuytren disease network with 13 potential risk factors based on the EBDMCMC approach. In Section 3.2, we consider only the severity measurements of the 10 fingers to infer the interaction among the

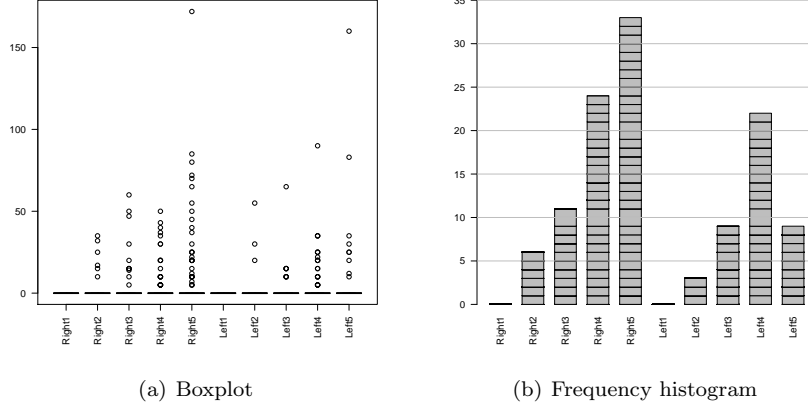


Figure 6: (a) Angles of the 10 fingers of all 279 patients. (b) Frequency histogram of rays affected with Dupuytren disease for all 10 fingers.

fingers.

3.1 Inference for Dupuytren disease with risk factors

We apply our Bayesian framework to infer the conditional (in)dependence among the 23 variables and to identify the potential risk factors of Dupuytren disease and discover how they affect the disease. In implementing the proposed EBDMCMC approach to analyze this dataset, we place a uniform distribution as an uninformative prior on the graph and the G-Wishart prior $W_G(3, I_{23})$ on the precision matrix. We run the EBDMCMC algorithm for 2,000K iterations with a 1,000K sweeps burn-in. The results are displayed in Figures 7 and 8.

The graph with the highest posterior probability is the graph with 42 edges. Figure 7 visualizes 26 edges that have posterior probabilities large than 0.5. Similarly, Figure 8 shows the image of all posterior inclusion probabilities where the degree of darkness increases with increasing posterior probabilities.

The edges in the graph show the interactions among the 10 severity measurements of Dupuytren disease and 13 risk factors. For example, the results (Figures 7 and 8) show that factors “Age”, “Alcohol”, “Ledderhose disease”, “Hand Injury” and “Relative”, among those 13 risk factors, have a significant association with the severity of Dupuytren disease. Graph 7 also shows that factor “Age” is a hub in this graph and it plays a significant role as it affects the severity of the disease directly and indirectly through the influence of other risk factors such as “Ledderhose”.

Further we checked the stability of the selected graph with highest posterior probability at convergence of the algorithm with 100 different starting points. The resulting Figure 9 shows the traces of number of edges in the estimated graphs plotted against iterations of the EBDMCMC algorithm with the 100 different starting points. The plot shows good mixing around a stable graph model size which is 42 and the algorithm converges after around 300 iterations.

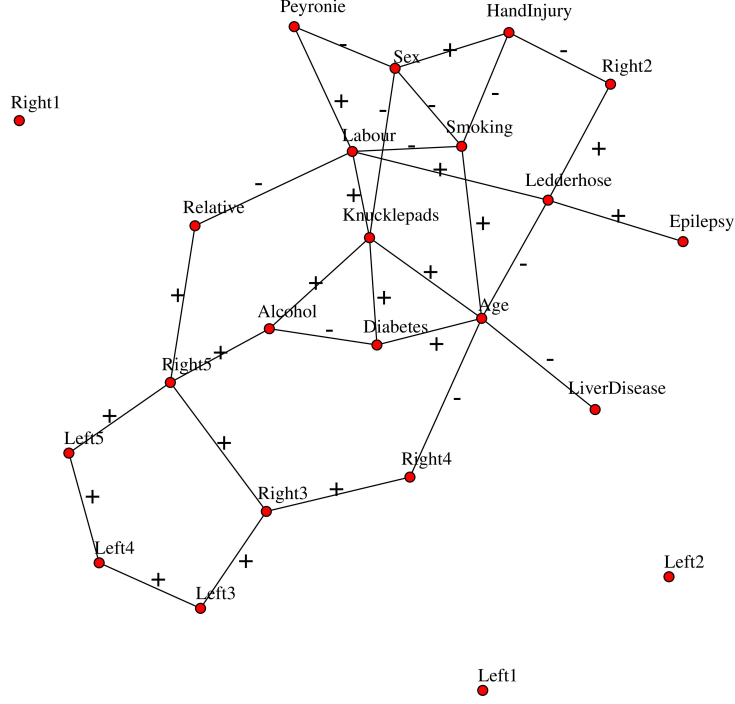


Figure 7: The inferred graph for the Dupuytren disease dataset based on 13 risk factors and the total degrees of flexion in all 10 fingers. It reports the selected graph with 26 edges for which their posterior inclusion probabilities (12) are more than 0.5. Sign “+” shows a positive relationship between nodes and “-” shows a negative relationship.

3.2 Severity of Dupuytren disease between pairs of fingers

Here, we consider the relationship between incurrence of Dupuytren disease in pairs of fingers on both right and left hands. Interaction between fingers is important because it help surgeons to decide whether they should operate one finger or multiple fingers simultaneously. The main idea is that if fingers are almost independent in terms of the severity of Dupuytren disease, there is no reason to operate the fingers simultaneously. On the other hand, if there is a strong relationship between fingers, then joint surgery may be recommended if one of the fingers is affected.

We apply the proposed EBDPMC approach for the 10 variables of Dupuytren disease severity measures using a uniform distribution as an uninformative prior on the graph and the G-Wishart $W_G(3, I_{10})$ prior on the precision matrix. We run the EBDPMC algorithm for 2,000K iterations with a 1,000K as burn-in. The results are displayed in Figure 10 and Figure 11.

The graph with the highest posterior probability is the graph with 12 edges. Figure 10 visualizes the selected graph with 8 edges, for which the posterior inclusion probabilities in (12) are greater than 0.5. The edges in the graph show the interactions among the fingers with regard to the severity of Dupuytren

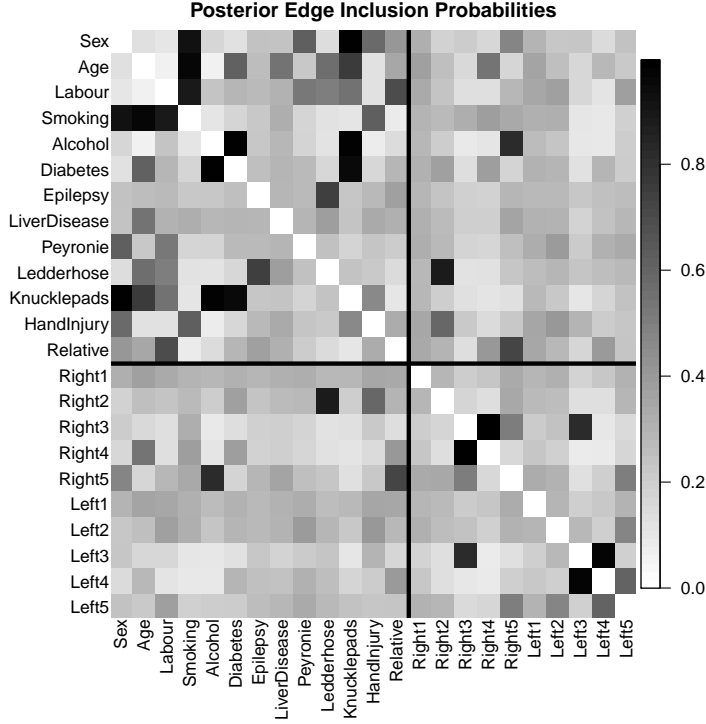


Figure 8: Image visualization of the posterior edge inclusion probabilities of all possible edges in the graph, for 10 fingers with 13 risk factors.

disease.

The results (Figures 10 and 11) show significant co-occurrences of Dupuytren disease in the ring fingers and middle fingers in both hands. This suggests that presence of disease in the middle finger is strongly associated to the ulnar side of the hand. Surprisingly, our results also show a strong relationship between middle fingers in both hands. Moreover, the results show that the joint interactions between fingers in both hands are almost symmetric. These results are in support of the hypotheses that the disease has genetic factors or other biological factors that affect similar fingers in both hands

3.3 Fit of model to Dupuytren data

Posterior predictive checks can be used for checking whether the proposed Bayesian approach fits the Dupuytren data set well or not. If the model fits the Dupuytren data, then simulated data generated under the model should look like to the observed data. In this regard, first, based on our estimated graph from the EBDMMCMC algorithm in section 3.1, we draw simulated data from the posterior predictive distribution. Then, we compare the samples to our observed data. Any systematic differences between the simulations and the data determine potential failings of the model.

We obtain the conditional distributions of the potential risk factors and disease severity measures on the fingers for both simulated and observed data.

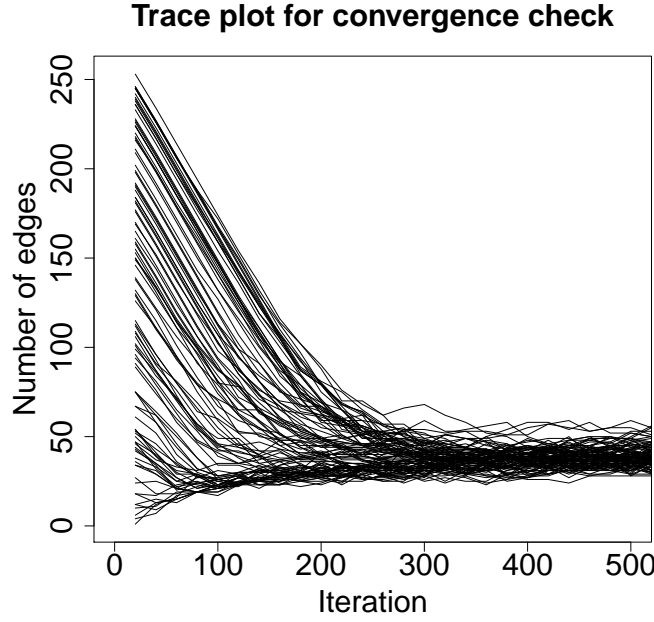


Figure 9: Case study of Section 3.1. Trace plot of the number of edges included in the estimated graphs against iterations of the EBDMCMC algorithm with 100 different starting points.

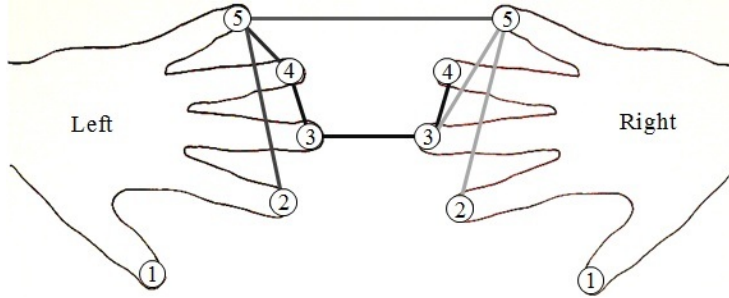


Figure 10: The inferred graph of Dupuytren disease dataset based on the total degrees of flexion in all 10 fingers. It reports the selected graph with 8 edges for which their posterior inclusion probabilities (12) are more than 0.5.

The empirical and predictive conditional distributions of some selected variables are presented in Figures 12, 13 and 14.

Figure 12 displays the empirical and predictive distributions of disease severity measure on finger 4 in right hand (right4) conditional on variable “age” in four categories $\{(40, 50), (50, 60), (60, 70), (70, 90)\}$. The variable right4, based on Tubiana Classification, grouped into 5 categories, category 1: 0 degree for total angle; 2: degree between (1, 45); 3: degree between (46, 90); 4: degree between (90, 135); 5: degree more than 135. The results in Figure 12 show that the fit is good, since the predicted conditional distributions, in general, are the same with the empirical distributions.

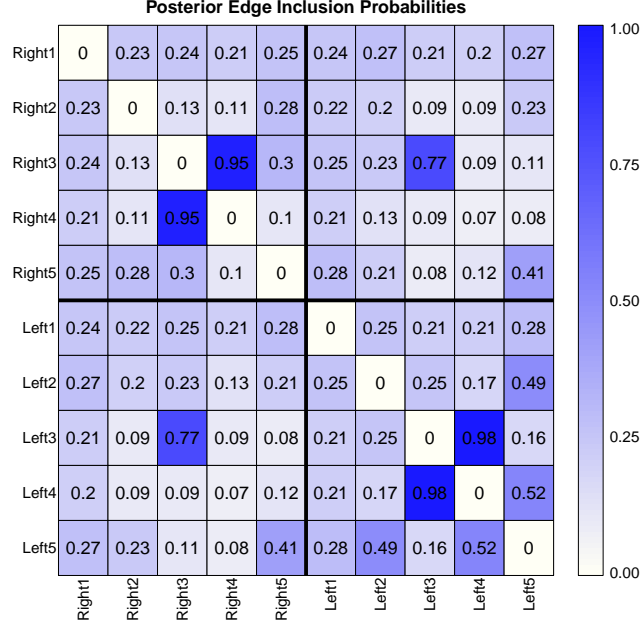


Figure 11: Image visualization of the posterior edge inclusion probabilities of all possible edges in the graph, for 10 fingers.

Similarly, Figure 13 plots the empirical and predictive distribution of disease severity measure on finger 5 in right hand (right5) conditional on variable “Relative” and Figure 14 plots the empirical and predictive distribution of disease severity measure on finger 2 in right hand (right2) conditional on variable “Ledderhose”. These results also suggest that the EBDMCMC algorithm fits the Dupuytren data well as the predicted conditional distributions are in agreement with the empirical distributions.

4 Conclusion

In this paper we have implemented a Bayesian method for discovering the effect of mixed potential risk factors of Dupuytren disease and the underlying relationships between fingers on both right and left hands with regard to severity of the Dupuytren disease.

The results of the case study clearly demonstrate that age, alcohol, relative, and ledderhose diseases all affect Dupuytren disease directly. Other risk factors only affect Dupuytren disease indirectly. Another important result is that severity of Dupuytren disease in fingers are correlated. In particular, the middle finger with the ring finger. This implies that a surgical intervention on either the ring or middle finger should preferably be executed simultaneously.

In Section 3, based on our Bayesian framework, we model the Dupuytren disease by consider the 13 potential risk factors. Our result in this section support the hypotheses the disease has genetic factors or other biological factors that affect the severity of the disease in fingers simultaneously. Indeed, in genome-wide association study, nine genes were identified to be associated with Dupuytren

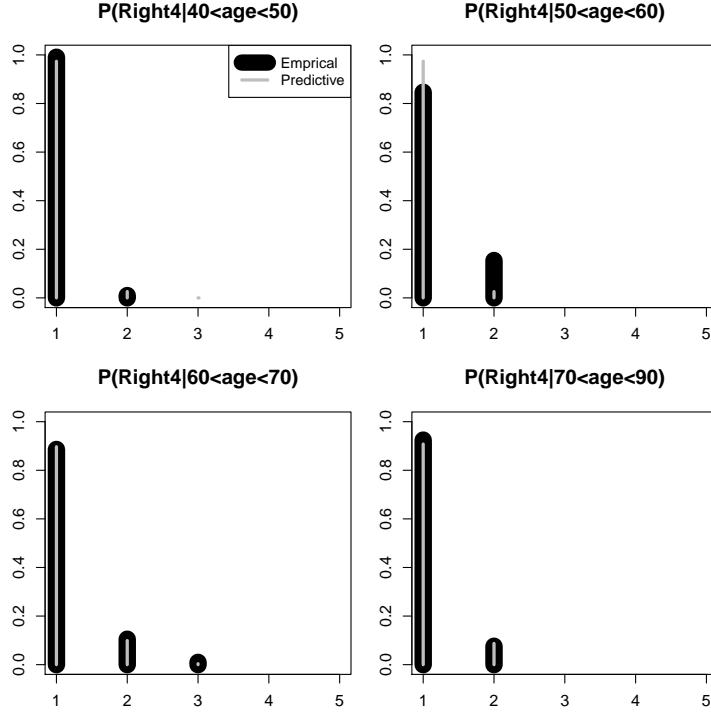


Figure 12: Empirical and predictive conditional distributions for total angle of finger 4 in right hand condition on four different categories of variable “age”.

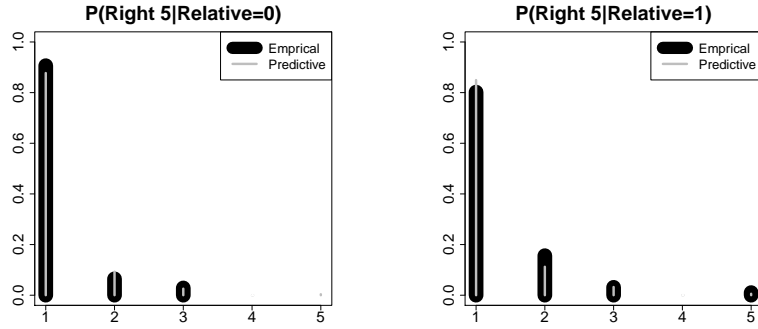


Figure 13: Empirical and predictive conditional distributions for total angles of finger 5 in right hand condition on relative variable.

disease (Dolmans et al., 2011). It should be interesting to consider the potential environmental risk factors jointly with those biological factors. Bayesian inference for all these factors requires a computationally efficient search algorithm that can potentially explore the underlying graph structure to uncover complicated patterns among these variables.

We compare our EBDMMCMC Bayesian approach with an alternative Bayesian approach (Dobra and Lenkoski, 2011) using a simulation study on various types of graph structures. Although, both approaches converge to the same posterior

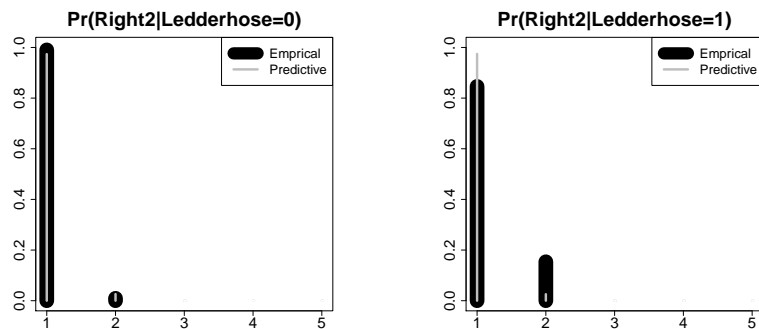


Figure 14: Empirical and predictive conditional distributions for total angles of finger 2 in right hand condition on Ledderhose disease variable.

distribution our approach has some clear advantages on finite MCMC runs. This difference is mainly due to our implementation of a computationally efficient algorithm. Our method is computationally more efficient because of two reasons. Firstly, our sampling algorithm is based on birth-death process. Secondly, we implemented an exact way of computing for the ratio of normalizing constants based on the result in Uhler et al. (2014), which has been computationally a bottle-neck in the Bayesian approach.

Of course, our extended Bayesian method is not limited only to this type of data. It can potentially be applied to any kind of mixed data where the observed variables are binary, ordinal or continuous.

References

- Albert, R. and A.-L. Barabási (2002). Statistical mechanics of complex networks. *Reviews of modern physics* 74(1), 47.
- Atay-Kayis, A. and H. Massam (2005). A monte carlo method for computing the marginal likelihood in nondecomposable gaussian graphical models. *Biometrika* 92(2), 317–335.
- Bayat, A. and D. McGrouther (2006). Management of dupuytren’s disease—clear advice for an elusive condition. *Annals of the Royal College of Surgeons of England* 88(1), 3.
- Cappé, O., C. Robert, and T. Rydén (2003). Reversible jump, birth-and-death and more general continuous time markov chain monte carlo samplers. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 65(3), 679–700.
- Cheng, Y., A. Lenkoski, et al. (2012). Hierarchical gaussian graphical models: Beyond reversible jump. *Electronic Journal of Statistics* 6, 2309–2331.
- Dempster, A. (1972). Covariance selection. *Biometrics* 28(1), 157–175.
- Dobra, A. and A. Lenkoski (2011). Copula gaussian graphical models and their application to modeling functional disability data. *The Annals of Applied Statistics* 5(2A), 969–993.

- Dolmans, G. H., P. M. Werker, H. C. Hennies, D. Furniss, E. A. Festen, L. Franke, K. Becker, P. van der Vlies, B. H. Wolffenbuttel, S. Tinschert, et al. (2011). Wnt signaling and dupuytren’s disease. *New England Journal of Medicine* 365(4), 307–317.
- Genest, C., K. Ghoudi, and L.-P. Rivest (1995). A semiparametric estimation procedure of dependence parameters in multivariate families of distributions. *Biometrika* 82(3), 543–552.
- Hoff, P. D. (2007). Extending the rank likelihood for semiparametric copula estimation. *The Annals of Applied Statistics*, 265–283.
- Jones, B., C. Carvalho, A. Dobra, C. Hans, C. Carter, and M. West (2005). Experiments in stochastic computation for high-dimensional graphical models. *Statistical Science* 20(4), 388–400.
- Lanting, R., D. C. Broekstra, P. M. Werker, and E. R. van den Heuvel (2014). A systematic review and meta-analysis on the prevalence of dupuytren disease in the general population of western countries. *Plastic and reconstructive surgery* 133(3), 593–603.
- Lanting, R., N. Noorae, P. Werker, and E. van den Heuvel (2014). Patterns of dupuytren disease in fingers; studying correlations with a multivariate ordinal logit model. *Plastic and reconstructive surgery*.
- Lanting, R., E. R. van den Heuvel, B. Westerink, and P. M. Werker (2013). Prevalence of dupuytren disease in the netherlands. *Plastic and reconstructive surgery* 132(2), 394–403.
- Lauritzen, S. (1996). *Graphical models*, Volume 17. Oxford University Press, USA.
- Lenkoski, A. (2013). A direct sampler for g-wishart variates. *Stat* 2(1), 119–128.
- Liang, F. (2010). A double metropolis–hastings sampler for spatial models with intractable normalizing constants. *Journal of Statistical Computation and Simulation* 80(9), 1007–1022.
- Meyerding, H. W. (1936). Dupuytren’s contracture. *Archives of Surgery* 32(2), 320–333.
- Milner, R. (2003). Dupuytren’s disease affecting the thumb and first web of the hand. *Journal of Hand Surgery (British and European Volume)* 28(1), 33–36.
- Mitsakakis, N., H. Massam, M. D. Escobar, et al. (2011). A metropolis-hastings based method for sampling from the g-wishart distribution in gaussian graphical models. *Electronic Journal of Statistics* 5, 18–30.
- Mohammadi, A. and E. Wit (2015a). *BDgraph: Bayesian Graph Selection Based on Birth-Death MCMC Approach*. R package version 2.22.
- Mohammadi, A. and E. C. Wit (2015b). Bayesian structure learning in sparse gaussian graphical models. *Bayesian Analysis* 10(1), 109–138.

- Murray, I., Z. Ghahramani, and D. MacKay (2012). Mcmc for doubly-intractable distributions. *arXiv preprint arXiv:1206.6848*.
- Powers, D. M. (2011). Evaluation: from precision, recall and f-measure to roc, informedness, markedness & correlation. *Journal of Machine Learning Technologies* 2(1), 37–63.
- Roverato, A. (2002). Hyper inverse wishart distribution for non-decomposable graphs and its application to bayesian inference for gaussian graphical models. *Scandinavian Journal of Statistics* 29(3), 391–411.
- Scutari, M. (2013). On the prior and posterior distributions used in graphical modelling. *Bayesian Analysis* 8(1), 1–28.
- Shih, B. and A. Bayat (2010). Scientific understanding and clinical management of dupuytren disease. *Nature Reviews Rheumatology* 6(12), 715–726.
- Sklar, M. (1959). *Fonctions de répartition à n dimensions et leurs marges*. Université Paris 8.
- Tubiana, R., B. Simmons, and H. DeFrenne (1982). Location of dupuytren’s disease on the radial aspect of the hand. *Clinical orthopaedics and related research* (168), 222.
- Uhler, C., A. Lenkoski, and D. Richards (2014). Exact formulas for the normalizing constants of wishart distributions for graphical models. *arXiv preprint arXiv:1406.4901*.
- Wang, H. (2012). Bayesian graphical lasso models and efficient posterior computation. *Bayesian Analysis* 7(4), 867–886.
- Wang, H. et al. (2015). Scaling it up: Stochastic search structure learning in graphical models. *Bayesian Analysis* 10(2), 351–377.
- Wang, H., C. M. Carvalho, et al. (2010). Simulation of hyper-inverse wishart distributions for non-decomposable graphs. *Electronic Journal of Statistics* 4, 1470–1475.
- Wang, H. and S. Li (2012). Efficient gaussian graphical model determination under g-wishart prior distributions. *Electronic Journal of Statistics* 6, 168–198.
- Wang, H. and N. S. Pillai (2013). On a class of shrinkage priors for covariance matrix estimation. *Journal of Computational and Graphical Statistics* 22(3), 689–707.
- Wong, F., C. K. Carter, and R. Kohn (2003). Efficient estimation of covariance selection models. *Biometrika* 90(4), 809–830.